

EV316937084

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

APPLICATION FOR LETTERS PATENT

Methods and Systems for Transparent Depth Sorting

Inventor(s):
Jeff Andrews

ATTORNEY'S DOCKET NO. ms1-1839us

RELATED APPLICATION

This application is a continuation-in-part of, and claims priority to U.S. Patent Application Serial No. 10/661,343, filed on September 12, 2003, the disclosure of which is incorporated by reference.

TECHNICAL FIELD

This invention relates to the field of computer graphics. More specifically, the present invention pertains to methods and systems for rendering objects.

BACKGROUND

Computer graphics systems typically utilize instructions, implemented via a graphics program on a computer system, to specify calculations and operations needed to produce two-dimensional or three-dimensional displays. Exemplary graphics systems that include APIs that are commercially available for rendering three-dimensional graphs include Direct3D, available from Microsoft Corporation, of Redmond, Wash., and OpenGL by Silicon Graphics, Inc., of Mountain View, Calif.

Computer graphics systems can be envisioned as a pipeline through which data pass, where the data are used to define an image that is to be produced and displayed. At various points along the pipeline, various calculations and operations that are specified by a graphics designer are used to operate upon and modify the data.

In the initial stages of the pipeline, the desired image is described by the application using geometric shapes such as lines and polygons, referred to in the art as "geometric primitives." The derivation of the vertices for an image and the

1 manipulation of the vertices to provide animation entail performing numerous
2 geometric calculations in order to eventually project the three-dimensional world
3 being synthesized to a position in the two-dimensional world of the display screen.

4 Primitives are constructed out of “fragments.” These fragments have
5 attributes calculated, such as color and depth. In order to enhance the quality of
6 the image, effects such as lighting, fog, and shading are added, and anti-aliasing
7 and blending functions are used to give the image a more realistic appearance.
8 The processes pertaining to per fragment calculation of colors, depth, texturing,
9 lighting, etc., are collectively known as “rasterization”.

10 The fragments and their associated attributes are stored in a frame buffer.
11 Once rasterization of the entire frame has been completed, pixel color values can
12 then be read from the frame buffer and used to draw images on the computer
13 screen.

14 To assist in understanding a typical computer graphics system, consider
15 Fig. 1 which illustrates, generally at 100, a system that can implement a computer
16 graphics process. System 100 comprises a graphics front end 102, a geometry
17 engine 104, a rasterization engine 106, and a frame buffer 108. System 100 can
18 typically be implemented in hardware, software, firmware or combinations
19 thereof, and is also referred to as a “rendering pipeline”.

20 Graphics front end 102 comprises, in this example, an application,
21 primitive data generation stage 102a and display list generation stage 102b. The
22 graphics front end generates geographic primitive data consumed by the
23 subsequent pipeline stage(s). Geographic primitive data is typically loaded from a
24 computer system’s memory and saved in a display list in the display list stage
25 102b. All geometric primitives are eventually described by vertices or points.

1 Geometry engine 104 comprises, in this example, high order surface (HOS)
2 tessellation 104a, and per-vertex operations stage 104b. In stage 104a, primitive
3 data is converted into simple rasterizer-supported primitives (typically triangles)
4 that represent the surfaces that are to be graphically displayed. Some vertex data
5 (for example, spatial coordinates) are transformed by four-by-four floating point
6 matrices to project the spatial coordinates from a position in the three-dimensional
7 world to a position on the display screen. In addition, certain other advanced
8 features can also be performed by this stage. Texturing coordinates may be
9 generated and transformed. Lighting calculations can be performed using the
10 vertex, the surface normal, material properties, and other light information to
11 produce a color value. Perspective division, which is used to make distant objects
12 appear smaller than closer objects in the display, can also occur in per-vertex
13 operations stage 104b.

14 Rasterization engine 106 is configured to perform so-called rasterization of
15 the re-assembled rasterizer-supported primitives. It comprises the following
16 stages: triangle/point assembly 106a, setup 106b, parametric evaluation 106c,
17 depth and stencil operations stage 106d, per-fragment operations stage 106e, and
18 the blend and raster operations (ROP) stage 106f.

19 Rasterization refers to the conversion of vertex data connected as rasterizer-
20 supported primitives into “fragments.” Each fragment corresponds to a single
21 element (e.g., a “pixel” or “sub-pixel”) in the graphics display, and typically
22 includes data defining color, transparency, depth, and texture(s). Thus, for a
23 single fragment, there are typically multiple pieces of data defining that fragment.
24 To perform its functions, triangle/point assembly stage 106a fetches different
25 vertex components, such as one or multiple texture component(s), a color

1 component, a depth component, and an alpha blending component (which
2 typically represents transparency).

3 Setup stage 106b converts the vertex data into parametric function
4 coefficients that can then be evaluated on a fragment coordinate (either pixel or
5 sub-pixel) by fragment coordinate basis. Parametric evaluation stage 106c
6 evaluates the parametric functions for all the fragments which lie within the given
7 rasterizable primitive, while conforming to rasterizable primitive inclusion rules
8 and contained within the frame buffer extents.

9 Depth and stencil operations stage 106d perform depth operations on the
10 projected fragment depth and application specified fragment stencil operations.
11 These operations apply to both the comparison function on the depth and stencil
12 values, how the depth and stencil values should be updated in the depth/stencil
13 buffer and whether the fragment should terminate or continue processing. In the
14 idealized rasterization pipeline these operations take place just before frame buffer
15 write-back (after blend and ROP stage 106f), but commonly these operations are
16 valid before the per-fragment operations stage 106e, which enables early
17 termination of many fragments and corresponding performance
18 optimizations/improvements.

19 Per-fragment operations stage 106e typically performs additional
20 operations that may be enabled to enhance the detail or lighting effects of the
21 fragments, such as texturing, bump mapping, per-fragment lighting, fogging, and
22 other like operations. Near the end of the rasterization pipeline is the blend and
23 raster operation (ROP) stage 106f, which implements blending for transparency
24 effects and traditional 2D blit raster operations. After completion of these
25 operations, the processing of the fragment is complete and it is typically written to

1 frame buffer 110 and potentially to the depth/stencil buffer 108. Thus, there are
2 typically multiple pieces of data defining each pixel.

3 Now consider so-called “depth sorting” as it pertains to rendering 3D
4 graphics. Depth sorting in computer graphics is typically accomplished using
5 what is referred to as a “depth buffer”. A depth buffer, often implemented as a “z-
6 buffer” or a “w-buffer”, is a 2D memory array used by a graphics device that
7 stores depth information to be accessed by the graphics device while rendering a
8 scene. Typically, when a graphics device renders a 3D scene to a render surface, it
9 can use the memory in an associated depth buffer surface as a workspace to
10 determine how the pixels or sub-pixels of rasterized polygons occlude one another.
11 The render surface typically comprises the surface or buffer to which final color
12 values are written. The depth buffer surface that is associated with the render
13 target surface is used to store depth information that tells the graphics device how
14 deep each visible pixel or sub-pixel is in the scene.

15 When a 3D scene is rasterized in a rasterization pipeline with depth
16 buffering enabled, each point on the rendering surface is typically tested. A depth
17 buffer that uses z values is often called a z-buffer, and one that uses w values is
18 called a w-buffer. Implementations may alternatively use other depth values such
19 as $1/z$ or $1/w$. While these invert the sense of increasing values with increasing
20 depth, this fact is typically hidden from the application and can therefore just be
21 thought of as a simple z or w buffer.

22 At the beginning of rendering a scene to a render target surface, the depth
23 value in the depth buffer is typically set to the largest possible value for the scene.
24 The color value on the rendering surface is set to either the background color value
25 or the color values of the background texture at that point. Once a fragment has

1 been generated at a given coordinate (x,y) on the rendering surface, the depth
2 value—which will be, for example, the z coordinate in a z-buffer, and the w
3 coordinate in a w-buffer—at the current coordinate is tested to see if it is smaller
4 than the depth value stored in the depth buffer. If the depth value of the polygon is
5 smaller, it is stored in the depth buffer and the color value from the polygon is
6 written to the current coordinate on the rendering surface. If the depth value of the
7 polygon at that coordinate is larger, the fragment is terminated, so the depth buffer
8 retains the smallest value at the current coordinate. This process is shown for
9 opaque polygons diagrammatically in Fig. 2.

10 There, notice that two polygons 200, 202 overlap along a ray that is
11 associated with a current coordinate 204 of interest. When the 3D scene is
12 rasterized, each coordinate on the rendering surface is typically tested. Here, the
13 corresponding location in depth buffer 206 corresponding to pixel (or sub-pixel)
14 204 is set to the largest possible value for the scene. The color value on the
15 rendering surface for this location can be set to a background color. Polygons 200
16 and 202 are effectively tested during rasterization to ascertain whether they
17 intersect with the current coordinate on the rendering surface. Since both
18 polygons intersect with the current coordinate on the rendering surface, the depth
19 value of polygon 200 at the current coordinate is effectively tested to see whether
20 it is smaller than the value at the current coordinate in the depth buffer. Here,
21 since the depth value of the polygon 200 for the associated coordinate is smaller
22 than the current depth value, the depth value for polygon 200 at the current
23 coordinate is written to the depth buffer and the color value for the polygon at the
24 current coordinate is written to the corresponding location in the rendering surface
25 (also referred to as the color buffer). Next, with the depth buffer holding the depth

1 value for polygon 200 at the current coordinate, the depth value of polygon 202 at
2 the current coordinate is tested against the current depth value in the depth buffer.
3 Since the depth value of polygon 202 at the current coordinate is smaller than the
4 depth value of polygon 200 at the current coordinate, the depth value of polygon
5 202 at the current coordinate is written to the corresponding depth buffer location
6 and the color value for polygon 202 at the current coordinate is written to the
7 corresponding location on the rendering surface.

8 In this manner, in the ultimately rendered image, overlapping portions of
9 polygon 202 at the current coordinate will occlude underlying portions of polygon
10 200 at the current coordinate.

11 Now consider what is a fundamental problem in 3D graphics—that which
12 is referred to as *transparent depth sorting*.

13 To appreciate this problem, consider that there are typically two different
14 types of pixels—opaque pixels and transparent pixels. Opaque pixels are those
15 pixels that pass no light from behind. Transparent pixels are those pixels that do
16 pass some degree of light from behind.

17 Consider now Fig. 3 which shows a viewer looking at a scene through one
18 exemplary pixel 300 on a screen. When an object is rendered by a 3D graphics
19 system, if the object is to appear as a realistic representation of what a viewer
20 would see in the real world, then this pixel should represent all of the light
21 contributions, reflected back towards the viewer, that lie along a ray R. In this
22 example, ray R intersects three different objects—an opaque mountain 302, a first
23 transparent object 304 and a second transparent object 306. The nearest opaque
24 pixel to the viewer is pixel 302a which lies on the mountain. Because this pixel is
25

1 opaque, no other pixels that might be disposed behind this pixel on the mountain
2 will make a contribution to the ultimately rendered pixel.

3 In the real world, transparent objects 304, 306 cause the light that is
4 reflected back towards the viewer to be affected in some way. That is, assume that
5 objects 304, 306 are glass or windows that have some type of coloration and
6 applied lighting and/or environmental effects. The effect of these windows is to
7 slightly dim or otherwise attenuate the light that is associated with pixel 302a then
8 apply lighting and/or environmental effects. In the real world, the viewer's eye
9 effectively sums all of the light contributions to provide a realistic image of the
10 distant mountain. In the 3D graphics world, this is not as easy.

11 Specifically, assume that pixel 302a has associated color values that
12 describe how that pixel is to be rendered without any transparency effects applied.
13 The influence of the right side of object 304 at 304a, and the left side of object 304
14 at 304b, with applied lighting 308, will change the color values of the associated
15 pixel. Similarly, the influence of the right side of object 306 at 306a, and the left
16 side at 306b, with applied lighting 310, will further change the color values of the
17 associated pixel.

18 Thus, if one wishes to accurately render pixel 302a, one should necessarily
19 take into account the transparency and lighting effects of these lighted transparent
20 objects, which requires back to front drawing ordering of the transparent objects'
21 contributions to the given pixel.

22 The traditional depth buffering techniques described above do nothing to
23 alleviate the back to front rendering order problem. Specifically, the traditional
24 depth buffering techniques essentially locate the closest pixel (i.e. the pixel with
25 the smallest z value) and then write the pixel's color values to the color buffer.

1 There is no back to front ordering, with partial application of the overlying pixel's
2 color values (and inversely partial retaining of the current color value). Thus,
3 traditional depth buffering techniques do not take into account this transparency
4 issue.

5 There have been attempts in the past to solve this particular transparency
6 depth sorting issue. For example, one solution to this problem is to push the
7 problem onto the application programmer. For example, the application
8 programmer might resolve this issue by drawing all of the opaque objects first, and
9 then perform some type of inexpensive bounding box or bounding sphere
10 processing, and present the resulting data to a graphics engine in back-to-front
11 order. This can unnecessarily burden the application programmer.

12 Another general scheme to attempt to solve the transparency depth sorting
13 problem is known as the "A-buffer" approach. This approach creates a per pixel
14 linked list of all of the pieces of per pixel data as a frame is being drawn. For
15 every pixel in the frame buffer, there is a linked list of fragments. These
16 fragments embody the contributions of the various objects at the given coordinate
17 that are in the scene. The A-buffer approach is a very general method that
18 essentially collects all of the linked list data for each pixel, and after all of linked
19 list data is collected, resolves the back to front issues, on a pixel by pixel basis,
20 after the scene is completely drawn. In the context of a resource-rich environment
21 where time is not a factor, this approach is acceptable as the software program
22 simply operates on the data as the data is provided to it.

23 One problem with the A-buffer approach, however, is most easily
24 appreciated in environments that are not necessarily resource-rich, and where time
25 is, in fact, a factor, e.g. the gaming environment where it is desirable to render real

1 time 3D graphics. With the A-buffer approach, the size of the linked lists and all
2 of the data in the linked lists can be quite large. Today, it is not economical to
3 have a frame buffer that is so large as to support the size of the linked lists. While
4 the results that are produced using the A-buffer approach are nice, the costs
5 associated with attaining such results are not appropriate for the real time
6 environment.

7 Accordingly, this invention arose out of concerns associated with providing
8 improved graphics systems and methods.

9 10 **SUMMARY**

11 Methods and systems for transparent depth sorting are described. In
12 accordance with one embodiment, multiple depth buffers are utilized to sort depth
13 data associated with multiple transparent pixels that overlie one another in back to
14 front order. The sorting of the depth data enables identification of an individual
15 transparent pixel that lies closest to an associated opaque pixel. With the closest
16 individual transparent pixel being identified, the transparency effect of the
17 identified pixel relative to the associated opaque pixel is computed. If additional
18 overlying transparent pixels remain, a next closest transparent pixel relative to the
19 opaque pixel is identified and, for the next closest pixel, the transparency effect is
20 computed relative to the transparency effect that was just computed. This process
21 iterates until there are no remaining overlying pixels in any transparent objects
22 being rasterized.

BRIEF DESCRIPTION OF THE DRAWINGS

Fig. 1 is a block diagram that illustrates one type of computer graphics system.

Fig. 2 is a diagram that illustrates a depth sorting process that utilizes a depth buffer.

Fig. 3 is a diagram that illustrates aspects of a computer graphics scenario in which multiple transparent objects are employed.

Fig. 4 illustrates an exemplary system that can be utilized to implement the various embodiments described herein.

Fig. 5 is a block diagram that illustrates a transparent depth sorting component in accordance with one embodiment.

Fig. 6 is a flow diagram and graphic that illustrates steps in a method in accordance with one embodiment.

DETAILED DESCRIPTION

Overview

Reference will now be made in detail to exemplary embodiments, examples of which are illustrated in the accompanying drawings. The described embodiments are not intended to limit application of the claimed subject matter to only the specific embodiments described. On the contrary, the claimed subject matter is intended to cover alternatives, modifications and equivalents, which may be included within the spirit and scope of various features of the described embodiments.

Furthermore, in the following detailed description, numerous specific details are set forth in order to provide a thorough understanding of the described

1 embodiments. It is quite possible, however, for the various embodiments to be
2 practiced without these specific details, but with details that are different, but still
3 within the spirit of the claimed subject matter. In some instances, well-known
4 methods, procedures, components, and circuits that are ancillary to, but support
5 the claimed embodiments have not been described in detail so as not to
6 unnecessarily obscure aspects of the embodiments that are described.

7 Some portions of the detailed descriptions which follow are presented in
8 terms of procedures, logic blocks, processing, and other symbolic representations
9 of operations on data bits within a computer memory or cache. These descriptions
10 and representations are the means used by those skilled in the data processing arts
11 to most effectively convey the substance of their work to others skilled in the art.
12 In the present application, a procedure, logic block, process, or the like, is
13 conceived to be a self-consistent sequence of steps or instructions leading to a
14 desired result. The steps are those requiring physical manipulations of physical
15 quantities. Usually, although not necessarily, these quantities take the form of
16 electrical or magnetic signals capable of being stored, transferred, combined,
17 compared, and otherwise manipulated in a computer system. It has proven
18 convenient at times, principally for reasons of common usage, to refer to these
19 signals as transactions, bits, values, elements, symbols, characters, fragments,
20 pixels, pixel data, or the like.

21 In the discussion that follows, terms such as “processing,” “operating,”
22 “calculating,” “determining,” “displaying,” or the like, refer to actions and
23 processes of a computer system or similar electronic computing device. The
24 computer system or similar electronic computing device manipulates and
25 transforms data represented as physical (electronic) quantities within the computer

1 system memories, registers or other such information storage, transmission or
2 display devices.

3 The embodiments described below pertain to a graphics subsystem that is
4 programmed or programmable to operate upon geometric primitive data that is to
5 be ultimately rendered to some type of display device. This graphics subsystem
6 can comprise an entire graphics engine, including a transform engine for geometry
7 calculations, a raster component comprising one or more of texture components,
8 specular components, fog components, and alpha blending components, and any
9 other components that can process application submitted and subsequent
10 intermediate graphics data. In some embodiments, the graphics subsystem can be
11 embodied in an integrated circuit component.

12 13 **Exemplary System**

14 Fig. 4 illustrates an exemplary system 400 that can be utilized to implement
15 one or more of the embodiments described below. This system is provided for
16 exemplary purposes only and is not intended to limit application of the claimed
17 subject matter.

18 System 400 exemplifies a computer-controlled, graphics system for
19 generating complex or three-dimensional images. Computer system 400
20 comprises a bus or other communication means 402 for communicating
21 information, and a processor 404 coupled with bus 402 for processing information.
22 Computer system 400 further comprises a random access memory (RAM) or other
23 dynamic storage device 406 (e.g. main memory) coupled to bus 402 for storing
24 information and instructions to be executed by processor 404. Main memory 406
25 also may be used for storing temporary variables or other intermediate information

1 during execution of instructions by processor 404. A data storage device 408 is
2 coupled to bus 402 and is used for storing information and instructions.
3 Furthermore, signal input/output (I/O) communication device 410 can be used to
4 couple computer system 400 onto, for example, a network.

5 Computer system 400 can also be coupled via bus 402 to an alphanumeric
6 input device 412, including alphanumeric and other keys, which is used for
7 communicating information and command selections to processor 404. Another
8 type of user input device is mouse 414 (or a like device such as a trackball or
9 cursor direction keys) which is used for communicating direction information and
10 command selections to processor 404, and for controlling cursor movement on a
11 display device 416. This input device typically has two degrees of freedom in two
12 axes, a first axis (e.g., x) and a second axis (e.g., y), which allows the device to
13 specify positions in a plane.

14 In accordance with the described embodiments, also coupled to bus 402 is a
15 graphics subsystem 418. Processor 404 provides graphics subsystem 418 with
16 graphics data such as drawing commands, coordinate vertex data, and other data
17 related to an object's geometric position, color, and surface parameters. In general,
18 graphics subsystem 418 processes the graphical data, converts the graphical data
19 into a screen coordinate system, generates pixel data (e.g., color, depth, texture)
20 based on the primitives (e.g., points, lines, polygons, and meshes), and performs
21 blending, anti-aliasing, and other functions. The resulting data are stored in a
22 frame buffer 420. A display subsystem (not specifically shown) reads the frame
23 buffer and displays the image on display device 416.

24 System 400 can be embodied as any type of computer system in which
25 graphics rendering is to take place. In some embodiments, system 400 can be

1 embodied as a game system in which 3D graphics rendering is to take place. As
2 will become apparent below, the inventive embodiments can provide more realistic
3 3D rendering in a real time manner that is particularly well suited for the dynamic
4 environments in which game system are typically employed.

6 Exemplary Embodiment Overview

7 In the embodiments described below, a solution to the transparent depth
8 sorting problem is provided. In accordance with the described embodiments,
9 multiple depth buffers are employed to support a sorting process that is directed to
10 incorporating the transparency effects of any intervening transparent pixels on an
11 associated opaque pixel that lies behind it.

12 In accordance with the described techniques, after all of the opaque pixels
13 are rendered to, for example, a rendering surface, the sorting process looks to
14 identify, for an associated opaque pixel, the closest transparent pixel that would
15 overlie it or, from a viewer's perspective, be in front of it. When the sorting
16 process finds the closest transparent pixel, the rasterization pipeline processes the
17 pixel data associated with the opaque pixel to take into account the effect that the
18 closest transparent pixel has on the opaque pixel. This typically results in the
19 color values for the pixel of interest being modified. Once the color values are
20 modified, the color values are written to the frame buffer for the pixel of interest.
21 The sorting process then looks to identify the next closest transparent pixel and, if
22 one is found, the rasterization pipeline processes the pixel data associated with the
23 current pixel of interest to take into account the effect that the next closest
24 transparent pixel has on the current pixel. Again, this typically results in the color
25 values for the pixel of interest being modified. Once the color values are

1 modified, the modified color values are again written to the frame buffer for the
2 pixel of interest. This sorting process continues until, for a given pixel, no
3 additional overlying transparent pixels are found.

4 The result of this process is that the transparency effects of overlying pixels
5 are taken into account, in back-to-front order, so that resultant pixels have
6 associated color values that accurately represent what the pixel would look like.

7 8 **Using Multiple Depth Buffers to Effect Transparent Depth Sorting**

9 In the example about to be described, multiple depth buffers are utilized to
10 effect transparent depth sorting. In this specific example, the depth buffers
11 comprise z-buffers. It is to be appreciated that the techniques described herein can
12 be employed in connection with w-buffers, $1/w$ buffers, $1/z$ buffers, or other depth
13 buffer types without departing from the spirit and scope of the claimed subject
14 matter. For example, in one embodiment an inverse w-buffer that is Huffman
15 encoded for varying precision throughout the range of depth can be used.

16 In the example about to be described, two physical z-buffers are employed.
17 The hardware in the graphics subsystem understands two pointers and parameters
18 for the two physical z-buffers. In accordance with one embodiment, both of the z-
19 buffers are readable, while only one of the z-buffers is writable. The z-buffer that
20 is both readable and writable is referred to as the "destination buffer", while the
21 readable only z-buffer is referred to as the "source buffer". Additionally, in the
22 embodiment about to be described, support is provided for the graphics software
23 to be able to flip which buffer is considered the source buffer and which buffer is
24 considered the destination buffer. The ability to flip which buffers are considered
25 as the source and destination buffers effectively provides for a logical-to-physical

1 mapping of the source and destination buffers to the actual physical buffers, as
2 will be appreciated by the skilled artisan. In one embodiment, flipping the buffers
3 is accomplished utilizing a multiplexer as input to the buffers address generation
4 function. To flip the buffers, the multiplexer select signal is flipped.

5 Fig. 5 shows a transparent depth sorting component 500 in accordance with
6 one embodiment. This component can reside at any suitable location within the
7 graphics processing pipeline. In one embodiment, component 500 can reside in
8 the graphics subsystem, such as subsystem 418 in Fig. 4. Additionally, while the
9 individual components of the transparent depth sorting component 500 are shown
10 to reside inside the component, such is not intended to constitute a physical
11 limitation on the arrangement and location of components. Rather, the transparent
12 depth sorting component can constitute a logical arrangement of components.
13 Accordingly, individual constituent parts of the transparent depth sorting
14 component 500 may not necessarily reside physically inside the component 500.

15 As illustrated, component 500 comprises or makes use of a first z buffer
16 502 (designated Z_0), a second z buffer 504 (designated Z_1), a z writeback counter
17 506 and comparison logic 508. A frame buffer 510 is provided and can be written
18 to as a result of the processing that is undertaken by transparent depth sorting
19 component 500, as will become apparent below.

20 Fig. 6 is a flow diagram that describes steps in a transparent depth sorting
21 method in accordance with one embodiment. The method can be implemented in
22 connection with any suitable hardware, software, firmware or combination thereof.
23 In one embodiment, the method can be implemented using a transparent depth
24 sorting component such as the one illustrated and described in Fig. 5. Alternately
25 or additionally, the method about to be described can be implemented using

1 various components within a graphics subsystem, such as the one shown and
2 described in Fig. 4.

3 In the method about to be described, two physical z-buffers are employed—
4 a first one designated z_0 and a second one designated z_1 . In the Fig. 6 flow
5 diagram, the contents in the individual buffers at various points in the described
6 process will be indicated, just to the right of the flow diagram, underneath
7 columns designated " z_0 " and " z_1 " which are grouped underneath encircled
8 numbers which indicate a particular pass in the process. So, for example, the first
9 pass in the process is designated by an encircled "1", the second pass in the
10 process is designated by an encircled "2" and so on. This is intended to help the
11 reader follow along the described process. Also notice just beneath the columns
12 designated " z_0 " and " z_1 " appears an example that shows, from a viewer's
13 perspective, four different pixels that lie along a particular ray. In this example,
14 pixel A is an opaque pixel, and pixels B, C, and D constitute transparent pixels (in
15 reverse order of their closeness to the viewer) which form the basis of the example
16 that is utilized to explain this embodiment. In this particular example, the depth
17 values of the pixels increase the further away from the viewer that they appear.
18 Thus, the depth value for pixel A is larger than the depth value for pixel B, and so
19 on.

20 In addition, recall that the notion of a "destination buffer" and a "source
21 buffer" was introduced above. The destination buffer is a buffer that is both
22 readable and writable, while the source buffer is a buffer that is only readable. In
23 the explanation that follows, the physical z buffers that are designated as the
24 destination and source buffers will change. To assist in keeping track of these
25

1 changes, a parenthetical "dest" and "src" will appear in the column associated with
2 the physical buffer that bears the associated designation.

3 Step 600 maps first z buffer z_0 as the destination z-buffer and step 602
4 renders all of the opaque objects. The mapping of the first z buffer operates to
5 designate the second z buffer z_1 as the source z buffer. Rendering the opaque
6 objects involves finding the opaque pixels that are the closest opaque pixels to the
7 viewer. Thus, rendering all of the opaque objects effectively obscures individual
8 pixels that lie behind the closest opaque pixel. When objects or pixels are
9 rendered, what is meant is that a series of primitives that represent the objects are
10 drawn. In some embodiments, this means that the whole series of triangles that
11 represent the objects are drawn.

12 When the opaque objects are rendered, individual color values for the
13 opaque objects are written to the frame buffer and their corresponding depth
14 values are written to the first z buffer z_0 . Thus, in this example, for the illustrated
15 ray, steps 600 and 602 result in the depth value for pixel A being written to the
16 first z buffer z_0 . At this point, the first z-buffer holds all of the depth data for all of
17 the opaque pixels associated with the rendering surface.

18 What the algorithm will now attempt to do, for a ray that is cast for a
19 particular x, y coordinate for the frame buffer, is find the transparent pixel
20 contribution that is as close as possible in depth and closer to the viewer than the
21 opaque pixel.

22 Accordingly, step 604 maps second z buffer z_1 as the destination z buffer
23 which effectively flips the logical mapping of the buffers, and steps 606 and 608
24 respectively clear the destination z buffer and initialize the destination z buffer to a
25 predetermined value. In this example, the predetermined value comprises its

1 smallest value. The smallest value in this example is the value associated with the
2 closest depth to the viewer—which is typically referred to as the “hither plane”.

3 Step 610 clears a “z writeback counter” (e.g. component 506 in Fig. 5),
4 which is utilized to keep track of writebacks that occur to the destination z buffer.
5 Step 612 sets the z compare logic to not write to the frame buffer, but to write to
6 the destination z buffer if a new z value is greater than the value in the destination
7 z buffer (i.e. the hither plane value) and the new z value is less than the value in
8 the source buffer (i.e. the depth value of pixel A). This step is directed to set up
9 the logic to ascertain the depth value of the transparent pixel that is closest and in
10 front of the opaque pixel. In this particular example, the result of performing step
11 614, which renders all of the transparent objects, is that the depth value associated
12 with pixel B is written into the destination buffer which, in this example, is the
13 second z buffer z_1 .

14 In the first pass the z buffer value is obtained. In the second pass (described
15 below), the color value for the z buffer level found in the first pass is obtained, and
16 the frame buffer is written to. This step involves drawing the series of primitives
17 that are associated with the transparent objects. Step 616 ascertains whether the z
18 writeback counter is equal to 0. If the z writeback counter is equal to 0, meaning
19 that no z writebacks occurred, then the method terminates or exits. Effectively,
20 exiting this process means that all of the transparency effects for the rendering
21 surface have been computed, back-to-front. In this particular single pixel
22 example, since a z writeback occurred to account for the depth value of pixel B,
23 the process does not exit.

24 Step 618 maps z_0 as the destination z buffer which effectively flips the
25 logical mapping of the z buffers. Step 620 sets the z compare logic to write to the

1 frame buffer (i.e. rendering surface) and the z buffer if the new z value is equal to
2 the value in the source z buffer (i.e. the value associated with pixel B) and the new
3 z value is less than the value in the destination z buffer (i.e. the depth value of
4 pixel A). Here, since the new z value is the value associated with pixel B, both
5 conditions are true, the frame buffer and the z buffer will be written to. Step 622
6 then renders all of the transparent objects, during which the value associated with
7 pixel B is written into the destination z buffer.

8 Accordingly, in the first pass, the process has found the backmost
9 transparent pixel (i.e. pixel B) and has written color values associated with the
10 transparency effects of pixel B out to the frame buffer. This provides, in the first
11 pass, a transparency effect that incorporates the contributions from pixel A and B.
12 In the pass about to be described, the transparency effects of pixel next closest to
13 pixel B (i.e. pixel C) will be computed. Following that pass, the transparency of
14 the next closest pixel (i.e. pixel D) will be computed.

15 Accordingly, step 622 returns to step 604 and maps second z buffer z_1 as
16 the destination z buffer which effectively flips the logical mapping of the buffers,
17 and steps 606 and 608 respectively clear the destination z buffer and initializes the
18 destination z buffer to its smallest value—i.e. the hither value. Step 610 clears the
19 z writeback counter and step 612 sets the z compare logic to not write to the frame
20 buffer, but to write to the destination z buffer if a new z value is greater than the
21 value in the destination z buffer (i.e. the hither plane value) and the new z value is
22 less than the value in the source buffer (i.e. the depth value of pixel B). This step
23 is directed to set up the logic to ascertain the closest transparent pixel to pixel B
24 which is also in front of pixel B. In this particular example, the result of
25 performing step 614, which renders all of the transparent objects, is that the depth

1 value associated with pixel C is written into the destination buffer which, in this
2 example, is the second z buffer z_1 .

3 Step 616 ascertains whether the z writeback counter is equal to 0. If the z
4 writeback counter is equal to 0, meaning that no z writebacks occurred, then the
5 method exits. Effectively, exiting this process means that all of the transparency
6 effects for the rendering surface have been computed, back-to-front. In this
7 particular single pixel example, since a z writeback occurred to account for the
8 depth value of pixel C, the process does not exit.

9 Step 618 maps z_0 as the destination z buffer which effectively flips the
10 logical mapping of the z buffers. Step 620 sets the z compare logic to write to the
11 frame buffer (i.e. rendering surface) and the z buffer if the new z value is equal to
12 the value in the source z buffer (i.e. the value associated with pixel C) and the new
13 z value is less than the value in the destination z buffer (i.e. the depth value of
14 pixel B). Here, since the new z value is the value associated with pixel C, both
15 conditions are true, the frame buffer and the z buffer will be written to. Step 622
16 then renders all of the transparent objects during which the value associated with
17 pixel C is written into the destination z buffer. At this point in the process, the
18 transparency effects due to pixel C have been accounted for in the rendering
19 surface.

20 The method then returns to step 604 and maps second z buffer z_1 as the
21 destination z buffer which effectively flips the logical mapping of the buffers, and
22 steps 606 and 608 respectively clear the destination z buffer and initializes the
23 destination z buffer to its smallest value—i.e. the hither value. Step 610 clears the
24 z writeback counter and step 612 sets the z compare logic to not write to the frame
25 buffer, but to write to the destination z buffer if a new z value is greater than the

1 value in the destination z buffer (i.e. the hither plane value) and the new z value is
2 less than the value in the source buffer (i.e. the depth value of pixel C). This step
3 is directed to set up the logic to ascertain the next closest transparent pixel to pixel
4 C that lies between pixel C and the viewer. In this particular example, the result of
5 performing step 614, which renders all of the transparent objects, is that the depth
6 value associated with pixel D is written into the destination buffer which, in this
7 example, is the second z buffer z_1 .

8 Step 616 ascertains whether the z writeback counter is equal to 0. If the z
9 writeback counter is equal to 0, meaning that no z writebacks occurred, then the
10 method exits. Effectively, exiting this process means that all of the transparency
11 effects for the rendering surface have been computed, back-to-front. In this
12 particular single pixel example, since a z writeback occurred to account for the
13 depth value of pixel D, the process does not exit.

14 Step 618 maps z_0 as the destination z buffer which effectively flips the
15 logical mapping of the z buffers. Step 620 sets the z compare logic to write to the
16 frame buffer (i.e. rendering surface) and the z buffer if the new z value is equal to
17 the value in the source z buffer (i.e. the value associated with pixel D) and the new
18 z value is less than the value in the destination z buffer (i.e. the depth value of
19 pixel C). Here, since the new z value is the value associated with pixel D, both
20 conditions are true, the frame buffer and the z buffer will be written to. Step 622
21 then renders all of the transparent objects, during which the value associated with
22 pixel D is written into the destination z buffer. At this point in the process, the
23 transparency effects due to pixel D have been accounted for in the rendering
24 surface. The method then returns to step 604.
25

1 Without going through the flow diagram again, with the depth value for
2 pixel D residing in the source z buffer due to the buffer flipping of step 604, the
3 flow diagram will effectively, at step 616, exit the process as the z writeback
4 counter will equal 0.

5 Thus, what has occurred at this point is that the individual transparent
6 pixels have been depth sorted, back-to-front, and their individual effects on given
7 pixels have been taken into account in a realistic manner. That is, the process
8 described above effectively sorts the z values of individual pixels to find the depth
9 of the transparent pixel that is furthest away from the viewer and closest to an
10 associated opaque pixel that the transparent pixel affects. The transparency effects
11 of this furthest most transparent pixel are computed and written to the frame
12 buffer. Then, the process repeats itself and looks for the transparent pixel that is
13 the next furthest most pixel away from the viewer and closest to the last
14 transparent pixel. If found, the transparency effects of the next furthest most
15 transparent pixel is computed and written to the frame buffer. This process
16 continues until there are no more transparent pixels along an associated ray. Thus,
17 the transparency effects of individual transparent pixels that lie along a particular
18 ray are taken into account in a back-to-front manner which provides realistic, real
19 time computer graphics rendering.

20 21 Conclusion

22 The methods and systems described above can provide solutions to the
23 transparent depth sorting process that can result in more realistic 3D graphics
24 rendering, particularly in scenarios in which real time rendering is appropriate.
25

1 Although the invention has been described in language specific to structural
2 features and/or methodological steps, it is to be understood that the invention
3 defined in the appended claims is not necessarily limited to the specific features or
4 steps described. Rather, the specific features and steps are disclosed as preferred
5 forms of implementing the claimed invention.

6

7

8

9

10

11

12

13

14

15

16

17

18

19

20

21

22

23

24

25